



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/694,923	10/27/2003	Michael M. Klock	NVID-062/001'S 140060-2128	5104
77306	7590	12/22/2008	EXAMINER	
NVIDIA C/O COOLEY GODWARD KRONISH LLP			WASHIBURN, DANIEL C	
Attn: Patent Group			ART UNIT	PAPER NUMBER
777 6th St NW				2628
Suite 1100				
WASHINGTON, DC 20001				
MAIL DATE		DELIVERY MODE		
12/22/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 10/694,923	Applicant(s) KLOCK ET AL.
	Examiner DANIEL WASHBURN	Art Unit 2628

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If no period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED. (35 U.S.C. § 133).

Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 17 October 2008.

2a) This action is FINAL. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1,21,25 and 28-30 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) Claim(s) _____ is/are allowed.

6) Claim(s) 1,21,25 and 28-30 is/are rejected.

7) Claim(s) _____ is/are objected to.

8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO/136/08)
Paper No(s)/Mail Date _____

4) Interview Summary (PTO-413)
Paper No(s)/Mail Date _____

5) Notice of Informal Patent Application
 6) Other: _____

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 10/17/08 has been entered.

Response to Arguments

Applicant's arguments with respect to the pending claims have been considered but are moot in view of the new ground(s) of rejection.

Claim Objections

Claim 25 is objected to because of the following informalities: the fifth and sixth lines from the bottom of claim 25 read, "...in response to detecting the **first threshold** being no greater than the second threshold level". For examination purposes, the examiner assumes it should read, "...in response to detecting the **percentage of stalled clock cycles** being no greater than the second threshold level", as this is along the same lines as the amendments made to the other independent claims and it does not create a conflict in claim language ("the first threshold level of stalls being greater than the second threshold level of stalls" ... "in response to detecting the first threshold being no greater than the second threshold level").

Appropriate correction is required.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1, 21, and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Giemborek et al. (US 6,950,105) in view of Williams et al. (US 6,397,343), in view of Oliver et al. (US 7,243,217) and further in view of Gulick (US 6,061,802).

As to claim 1, Giemborek describes a method of operating a graphics system having a sequence of at least two discrete performance levels with each performance level being defined by a core clock rate of a graphics processing unit and a memory clock rate, the method comprising:

operating the graphics system at the core clock rate and memory clock rate associated with a selected performance level, the selected performance level being a minimum performance level sufficient to maintain the display rate within a normal range (column 1 lines 50-67 through column 2 lines 1-11 and Figure 1 describes graphics accelerator 10, which matches the speed of at least one of two or more clocks (e.g., engine clock 40 and memory clock 42) to levels (speeds) under software control to a rate sufficient to satisfy current display requirements. The graphics accelerator includes 2D/3D engine 20, overlay engine 22, and frame buffer 16. Further, column 7 lines 50-67 through column 8 lines 1-27 describes increasing the clock speeds of the clocks within the graphics accelerator 10 if the system is currently over-utilized and decreasing

the clock speeds of the clocks within the graphics accelerator 10 if the system is currently under-utilized. Finally, column 1 lines 50-65 describes that the clock speeds are set such that display rate remains in a normal range (i.e., graphics display performance is not sacrificed)).

Giemborek doesn't describe monitoring utilization of a graphics pipeline.

However, Williams describes monitoring the utilization of a graphics pipeline.

Williams describes a method and system that includes a device for dynamic graphics subsystem clock adjustment within a computer system having a CPU and a dedicated graphics subsystem. A system interface is coupled to the graphics subsystem to allow a controller to determine the graphics processing load placed on the graphics subsystem (column 4 lines 11-31 and column 6 lines 33-50). Williams further describes that monitoring said at least one attribute (in this case the graphics processing load placed on the graphics subsystem) comprises: monitoring at least one attribute indicative of utilization of a graphics pipeline within a graphics processor core clock domain and determining whether the graphic pipeline is under-utilized or over-utilized (column 6 lines 51-67, column 7 lines 1-20, and column 8 lines 11-22 describes that the device 100 determines the graphics subsystem load by monitoring the processing activity of graphics subsystem 200 via the pipeline control 206 (e.g., by snooping graphics commands and data flowing through a graphics pipeline to determine the activity level of the graphics pipeline) and adjusts the pipeline clock frequency accordingly).

All the above-described limitations of claim 1 are known in Giemborek and Williams, the only difference is the combination of old elements into a single system and method.

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to include in Giemborek the system and method of determining the processing load placed on a graphics subsystem by monitoring the level of activity of a graphics pipeline, and adjusting the frequency of the pipeline clock according to the determined load, as taught by Williams, as this doesn't change the operation of the rest of the system, and it could be used to achieve the predictable results of improving the efficiency of the 2D/3D graphics engine disclosed in Giemborek by monitoring the pipeline activity within the graphics engine and determining a required clock rate to be passed to the engine based, in part, on the monitored activity. One advantage of passing the graphics engines a variable clock rate based at least in part on calculated activity within a pipeline within the engine is that the system can further reduce its power consumption and optimize its calculated clock rates by adding the activity of the pipeline to the list of factors that are used when determining the clock rates.

Giemborek in view of Williams doesn't describe monitoring a percentage of clock cycles at one or more points within a graphics processor core clock domain for which one or more stages of a graphics pipeline are stalled waiting for inputs from upstream stages as an indicator of utilization and determining whether the graphic pipeline is under-utilized or over-utilized based on comparing the percentage of clock cycles for which there is a stall within one or more blocks the graphics pipeline against a first

threshold level of stalls and a second threshold level of stalls with the first threshold level of stalls being greater than the second threshold level of stalls.

However, Oliver describes a system and method including monitoring a percentage of clock cycles at one or more points within a processor core clock domain for which one or more stages of a pipeline are stalled waiting for inputs from upstream stages as an indicator of utilization and determining whether the pipeline is under-utilized or over-utilized based on comparing the percentage of clock cycles for which there is a stall within one or more blocks the pipeline against a first threshold level of stalls (4:65-5:15 "The present invention decouples the clock speed of integer unit 110 and FPU 120 using command and data queues (or reservation stations) in dispatch unit 123 and control logic in execution pipeline clock controller 205. Execution pipeline clock controller 205 set the clock speed of FPU execution pipeline 124 as a function of the number and type of commands in the reservation stations in dispatch unit 123. This information is determined from Reservation Stations Full Levels status signals received from dispatch unit 123 and an Integer Pipe Stall Instruction signal received from issue unit 122. Execution pipeline clock controller 205 sets the speed of the Output Clock signal to a high rate (Fast mode) if the reservation stations are filling up, if integer unit 110 is stalled waiting for the result from FPU 120, or if the commands in the reservation stations require multiple cycles to execute." Further, 5:40-62 discloses "Execution pipeline clock controller 205 increases the Output Clock signal speed as the level rises in each reservation station or if an opcode indicates that integer unit 110 is waiting for a result from FPU 120 (process step 420). Execution pipeline clock controller 205 also

decreases the Output Clock signal speed as the level drops in each reservation station and if no queued opcode indicates that integer unit 110 is waiting for a result from FPU 120 (process step 425)." Thus, when the percentage of clock cycles at one or more points within the processor core clock domain for which the integer unit stage of the pipeline is stalled waiting for inputs from the upstream floating point unit stage is greater than zero and the floating point unit is being clocked at the slower speed this acts as an indicator in determining that the pipeline is under-utilized (at which point an opcode is sent to the floating point unit indicating that its clock speed should be increased), and when the percentage of clock cycles at one or more points within the processor core clock domain for which the integer unit stage of the pipeline is stalled waiting for inputs from the upstream floating point unit stage is not greater than zero and the floating point unit is being clocked at the faster speed this acts as an indicator in determining that the pipeline is under-utilized (at which point the clock speed of the floating point unit is reduced), where the determination is based on comparing the percentage of clock cycles for which there is a stall within one or more blocks the pipeline against a first threshold level of stalls.

Further, increasing the clock rate of the floating point unit as needed is considered increasing the performance level in response to detecting an over-utilization condition corresponding to the percentage of stalled clock cycles exceeding the first threshold level in order to increase the clock rate in the processor core clock domain, and decreasing the clock rate of the floating point unit as needed is considered decreasing the performance level in response to detecting an under-utilization condition

corresponding to the percentage of stalled clock cycles being no greater than the first threshold level to decrease the clock rate in the processor core clock domain.

All the elements above-described of claim 1 are known in Giemborek in view of Williams and further in view of Oliver, the only difference is the combination of known elements into a single system and method.

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to include in Giemborek in view of Williams the system and method of monitoring a percentage of clock cycles at one or more points within a graphics processor core clock domain for which one or more stages of a graphics pipeline are stalled waiting for inputs from upstream stages as an indicator of utilization and determining whether the graphic pipeline is under-utilized or over-utilized based on comparing the percentage of clock cycles for which there is a stall within one or more blocks of the graphics pipeline against a first threshold level of stalls; and increasing the performance level in response to detecting an over-utilization condition corresponding to the percentage of stalled clock cycles exceeding the first threshold level in order to increase the clock rate in the graphics processor core clock domain and decreasing the performance level in response to detecting an under-utilization condition corresponding to the percentage of clock cycles being no greater than the first threshold level to decrease the clock rate in the graphics processor core clock domain, as suggested by Oliver, in order to achieve the predictable result of improving the efficiency and power conservation of the 2D/3D graphics engine (disclosed in Giemborek) by monitoring the pipeline activity within the graphics engine and determining an appropriate clock rate.

The advantage of an adjustable clock rate is it allows the pipeline to meet processing demands when the required processing load is high while also conserving power when the required processing load is low.

Giemborek in view of Williams and further in view of Oliver doesn't describe the clock rate is adjusted according to two different threshold values, wherein the first threshold value is greater than the second threshold value, and further doesn't describe increasing the clock rate when the first threshold value is exceeded and decreasing the clock rate when the second threshold value is not exceeded.

However, Gulick describes a system and method wherein the clock rate is adjusted according to two different threshold values, wherein the first threshold value is greater than the second threshold value, and further describes increasing the clock rate when the first threshold value is exceeded and decreasing the clock rate when the second threshold value is not exceeded (8:62-9:5 "It is noted that in an alternative embodiment a plurality of threshold flags are used. A first threshold flag may indicate an upper bound of the desired clock rate and a second threshold flag may indicate a lower bound of the desired clock rate. If the data in the buffer causes the first threshold flag to be asserted, the frame clock rate is decreased. If the data in the buffer causes the second threshold flag to be asserted, the rate of the frame clock is increased. If the data in the data buffer remains at a level between the threshold flags, the clock adjustment is successful and no adjustment is necessary.").

All the elements of claim 1 are known in Giemborek, Williams, Oliver, and Gulick, the only difference is the combination of known elements into a single system and method.

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to include in Giemborek, Williams, and Oliver the system and method wherein the clock rate is adjusted according to two different threshold values, wherein the first threshold value is greater than the second threshold value, and further comprising increasing the clock rate when the first threshold value is exceeded and decreasing the clock rate when the second threshold value is not exceeded, as taught by Gulick, as this doesn't change the overall operation of the system, and it could be used to achieve the predictable result of creating a small percentage range of an acceptable amount of stall cycles in the system, which has the advantage of reducing power consumption, as the clock speed will not be increased for very brief stall periods, while still not affecting perceivable system performance, as a user will not be aware of the minimal and very brief delay that occurs within the pipeline.

Concerning claim 21, Giemborek describes a method of operating a graphics system having a sequence of at least two discrete performance levels where each performance level is defined by a core clock rate of a graphics processing unit and a memory clock rate, the performance levels including a high performance level for processing complex three-dimensional graphical images and at least one lower power, lower performance level for processing less complex graphical images, the method comprising:

operating the graphics system at the core clock rate and memory clock rate associated with the selected performance level, the selected performance level being a minimum performance level sufficient to maintain the display rate within the normal range (see the corresponding section in the rejection of claim 1).

Giemborek doesn't describe but Williams describes monitoring the utilization of a graphics pipeline (see the corresponding section in the rejection of claim 1).

Giemborek in view of Williams doesn't describe but Oliver describes monitoring a percentage of clock cycles at one or more points within a graphics processor core clock domain for which one or more stages of a graphics pipeline are stalled waiting for inputs from upstream stages as an indicator of utilization and determining whether the graphic pipeline is under-utilized or over-utilized based on comparing the percentage of clock cycles for which there is a stall within the graphics pipeline against a first threshold level of stalls (see the corresponding section in the rejection of claim 1);

in response to detecting a level of utilization greater than an over-utilization threshold for which a display rate of the graphics system is likely to be significantly decreased below a normal display rate, selecting a higher performance level to increase a clock rate in the graphics processor core clock domain, the over-utilization threshold corresponding to the percentage of stalled clock cycles exceeding the first threshold level (4:65-5:62, the variable speed clock applied to the floating point unit can easily be applied to one or more stages in a graphics processing pipeline (see 4:29-35), such as the graphics processing pipeline described in Williams. Thus, the combination is considered to suggest detecting a level of utilization greater than an over-utilization

threshold for which a display rate of the graphics system is likely to be significantly decreased below a normal display rate); and

in response to detecting a level of utilization below an under-utilization threshold, selecting a lower performance level to reduce the clock rate in the graphics processor core clock domain to reduce power required by the graphics system, the under-utilization threshold corresponding to the percentage of stalled clock cycles being no greater than the first threshold level (4:65-5:62, the variable speed clock applied to the floating point unit can easily be applied to one or more stages in a graphics processing pipeline (see 4:29-35), such as the graphics processing pipeline described in Williams).

See the rejection of claim 1 for rationale to combine Oliver with Giemborek and Williams.

Giemborek in view of Williams and further in view of Oliver doesn't describe but Gulick describes a system and method wherein the clock rate is adjusted according to two different threshold values, wherein the first threshold value is greater than the second threshold value, and further doesn't describe increasing the clock rate when the first threshold value is exceeded and decreasing the clock rate when the second threshold value is not exceeded. See the rejection of claim 1 for cited portions of Gulick and rationale.

As to claim 25, Giemborek describes a graphics system, comprising:
a graphics processor having a sequence of at least two discrete performance levels where each performance level is defined by a graphics processor core clock rate of a graphics processing unit and a memory clock rate (column 1 lines 50-67 through

column 2 lines 1-11 and Figure 1 describes graphics accelerator 10, which matches the speed of at least one of two or more clocks (e.g., engine clock 40 and memory clock 42) to levels (speeds) under software control to a rate sufficient to satisfy current display requirements. The graphics accelerator includes 2D/3D engine 20, overlay engine 22, and frame buffer 16); and

a graphics memory coupled to said graphics processor by a graphics bus and operable at said memory clock rate (Figure 1 and column 2 lines 27-36 describes frame buffer memory 16);

the graphics system operating at the core clock rate and memory clock rate associated with the performance level selected by the performance level controller, the selected performance level being a minimum performance level capable of maintaining the display rate within a normal range (column 1 lines 50-65 describes that the clock speeds are selected to ensure that the graphics display performance is not degraded).

Giemborek doesn't describe but Williams describes monitoring the utilization of a graphics pipeline (see the corresponding section in the rejection of claim 1).

Giemborek in view of Williams doesn't describe but Oliver describes a performance level controller, said performance level controller configured to monitor, as a function of time a percentage of clock cycles at one or more points within a processor core clock domain for which one or more stages of a pipeline are stalled waiting for inputs from upstream stages as an indicator of utilization and determining whether the pipeline is under-utilized or over-utilized by comparing the percentage of clock cycles for which there is a stall in the pipeline against a first threshold level of stalls (4:65-5:22 and

5:53-62 describes execution pipeline clock controller 205. Also see the corresponding section in the rejection of claim 1); and

 said performance level controller configure to increase said performance level to increase a clock rate in the processor core clock domain and to avoid over-utilization of said pipeline in response to detecting the percentage of stalled clock cycles exceeding the first threshold level (5:53-56).

 said performance level controller configured to decrease said performance level from a high performance level to a lower performance level to decrease the clock rate in the processor core clock domain to avoid under-utilization of said pipeline in response to detecting the percentage of stalled clock cycles being no greater than the second threshold value (5:57-62).

 See the rejection of claim 1 for rationale to combine Oliver with Giemborek and Williams.

 Giemborek in view of Williams and further in view of Oliver doesn't describe but Gulick describes a system and method wherein the clock rate is adjusted according to two different threshold values, wherein the first threshold value is greater than the second threshold value, and further doesn't describe increasing the clock rate when the first threshold value is exceeded and decreasing the clock rate when the second threshold value is not exceeded. See the rejection of claim 1 for cited portions of Gulick and rationale.

 Claims 28-30 are rejected under 35 U.S.C. 103(a) as being unpatentable over Giemborek et al. (US 6,950,105) in view of Williams et al. (US 6,397,343) in view of

Oliver et al. (US 7,243,217) and further in view of Gulick (US 6,061,802), as applied to claims 1, 21, and 25 above, and further in view of Culbert et al. (US 6,820,209).

As to claims 28-30, Giemborek describes a 2D/3D graphics engine that is capable of operating at different clock speeds (column 2 lines 42-67), where the 2D/3D graphics engine is operated at a speed that is determined based on factors that include the software running on a host CPU, as well as display mode settings of the computer, such as screen resolution, pixel or color depth, and screen refresh rate (column 5 lines 13-48).

Giemborek in view of Williams and further in view of Oliver and Gulick doesn't describe the method of claims 1, 21, or 25 wherein said at least two discrete performance levels include a low power two-dimensional graphics performance level, a standard two-dimensional graphics performance level, a low power three-dimensional graphics performance level, and a high performance three-dimensional graphics performance level.

However, Culbert describes a system and method wherein at least two discrete performance levels include a low power two-dimensional graphics performance level, a standard two-dimensional graphics performance level, a low power three-dimensional graphics performance level, and a high performance three-dimensional graphics performance level (column 5 lines 46-67 through column 6 lines 1-43 describes a system that includes a 2D graphics engine 212 and a separate 3D graphics engine 214. The 2D and 3D graphics engines are only activated when the graphics controller needs to produce 2D or 3D graphics. Further, column 6 lines 44-67 through column 7 lines 1-

16 describes that the graphics controller supplies a clock signal to the 2D engine and a second clock signal to the 3D engine. To reduce power consumption, the clock signal normally sent to the 2D engine is stopped when the 2D engine is not being utilized, and likewise the clock signal normally sent to the 3D engine is stopped when its processing resources are not being utilized. Thus, Culbert is considered to describe a low power two-dimensional graphics performance level (when the 2D processor isn't being utilized), a standard two-dimensional graphics performance level (when the 2D processor is being utilized), a low-power three-dimensional graphics performance level (when the 3D processor isn't being utilized), and a high performance three-dimensional graphics performance level (when the 3D processor is being utilized)).

All the elements of claims 28-30 are known in Giemborek, Williams, Oliver, Gulick, and Culbert, the only difference is the combination of known elements into a single system and method.

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to include in Giemborek in view of Williams and further in view of Oliver and Gulick the system and method wherein said at least two discrete performance levels include a low power two-dimensional graphics performance level, a standard two-dimensional graphics performance level, a low power three-dimensional graphics performance level, and a high performance three-dimensional graphics performance level, as taught by Culbert, as a system and method wherein the 2D engine is separate from the 3D engine, and where each engine is separately controlled, could be used to achieve the predictable result of more effectively saving power, as when only the 2D or

only the 3D engine is required for a particular operation, the other engine can be powered down, which doesn't degrade the performance of the system, and reduces power consumption even more than clocking the 2D and 3D engines up and down together (as described in Giemborek).

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to DANIEL WASHBURN whose telephone number is (571)272-5551. The examiner can normally be reached on Monday through Friday 8:30 a.m. to 5:00 p.m..

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Ulka Chauhan can be reached on (571) 272-7782. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Dan Washburn/
Examiner, Art Unit 2628
12/17/08

/Ulka Chauhan/
Supervisory Patent Examiner, Art
Unit 2628